

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



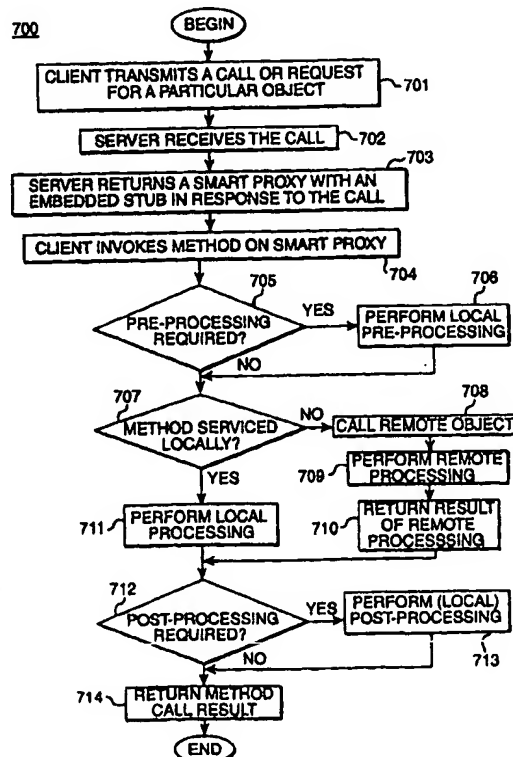
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/46		A1	(11) International Publication Number: WO 99/44123 (43) International Publication Date: 2 September 1999 (02.09.99)
(21) International Application Number: PCT/US99/03221 (22) International Filing Date: 16 February 1999 (16.02.99) (30) Priority Data: 60/076,048 26 February 1998 (26.02.98) US 09/044,930 20 March 1998 (20.03.98) US (71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, MS UPAL01-521, Palo Alto, CA 94303 (US). (72) Inventors: ARNOLD, Kenneth, C., R., C.; 7 Moon Hill Road, Lexington, MA 02173 (US). WALDO, James, H.; 155 Ruby Road, Dracut, MA 01826 (US). SCHEIFLER, Robert; 96 North Street, Somerville, MA 02144 (US). WOLLRATH, Ann, M.; 9 Northwoods Road, Groton, MA 01450 (US). (74) Agents: GARRETT, Arthur, S.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US) et al.		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: DOWNLOADABLE SMART PROXIES FOR PERFORMING PROCESSING ASSOCIATED WITH A REMOTE PROCEDURE CALL IN A DISTRIBUTED SYSTEM

(57) Abstract

Use of a smart proxy as a wrapper around a stub in a distributed system. Instead of receiving a stub as a result of a remote procedure call, a caller receives a smart proxy including the stub as an embedded object. The smart proxy performs predefined processing associated with a remote procedure call, the processing possibly occurring before, during, or after a response to the call.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

**DOWNLOADABLE SMART PROXIES FOR PERFORMING
PROCESSING ASSOCIATED WITH A REMOTE
PROCEDURE CALL IN A DISTRIBUTED SYSTEM**

REFERENCE TO RELATED APPLICATIONS

The following identified U.S. patent applications are relied upon and are incorporated by reference in this application as if fully set forth.

Provisional U.S. Patent Application No. 60/076,048, entitled "Distributed Computing System," filed on February 26, 1998.

U.S. Patent Application No. 09/044,923, entitled "Method and System for Leasing Storage," bearing attorney docket no. 06502.0011-01000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,838, entitled "Method, Apparatus, and Product for Leasing of Delegation Certificates in a Distributed System," bearing attorney docket no. 06502.0011-02000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,834, entitled "Method, Apparatus and Product for Leasing of Group Membership in a Distributed System," bearing attorney docket no. 06502.0011-03000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,916, entitled "Leasing for Failure Detection," bearing attorney docket no. 06502.0011-04000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,933, entitled "Method for Transporting Behavior in Event Based System," bearing attorney docket no. 06502.0054-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,919, entitled "Deferred Reconstruction of Objects and Remote Loading for Event Notification in a Distributed System," bearing attorney docket no. 06502.0062-01000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,938, entitled "Methods and Apparatus for Remote Method Invocation," bearing attorney docket no. 06502.0102-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/045,652, entitled "Method and System for Deterministic Hashes to Identify Remote Methods," bearing attorney docket no. 06502.0103-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,790, entitled "Method and Apparatus for Determining Status of Remote Objects in a Distributed System," bearing attorney docket no. 06502.0104-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,917, entitled "Suspension and Continuation of Remote Methods," bearing attorney docket no. 06502.0106-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,835, entitled "Method and System for Multi-Entry and Multi-Template Matching in a Database," bearing attorney docket no. 06502.0107-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,839, entitled "Method and System for In-Place Modifications in a Database," bearing attorney docket no. 06502.0108, and filed on the same date herewith.

U.S. Patent Application No. 09/044,945, entitled "Method and System for Typesafe Attribute Matching in a Database," bearing attorney docket no. 06502.0109-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,931, entitled "Dynamic Lookup Service in a Distributed System," bearing attorney docket no. 06502.0110-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,939, entitled "Apparatus and Method for Providing Downloadable Code for Use in Communicating with a Device in a Distributed System," bearing attorney docket no. 06502.0112-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,826, entitled "Method and System for Facilitating Access to a Lookup Service," bearing attorney docket no. 06502.0113-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,932, entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," bearing attorney docket no. 06502.0114-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/030,840, entitled "Method and Apparatus for Dynamic Distributed Computing Over a Network," bearing attorney docket no. 06502.0115-00000, and filed on February 26, 1998.

U.S. Patent Application No. 09/044,936, entitled "An Interactive Design Tool for Persistent Shared Memory Spaces," bearing attorney docket no. 06502.0116-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,934, entitled "Polymorphic Token-Based Control," bearing attorney docket no. 06502.0117-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,915, entitled "Stack-Based Access Control," bearing attorney docket no. 06502.0118-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,944, entitled "Stack-Based Security Requirements," bearing attorney docket no. 06502.0119-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,837, entitled "Per-Method Designation of Security Requirements," bearing attorney docket no. 06502.0120-00000, and filed on the same date herewith.

FIELD OF THE INVENTION

The present invention relates to a system and method for transmitting objects between machines in a distributed system and more particularly relates to transmission of a representation of a remote object including code for local processing.

BACKGROUND OF THE INVENTION

Distributed programs which concentrate on point-to-point data transmission can often be adequately and efficiently handled using special-purpose protocols for remote terminal access and file transfer. Such protocols are tailored specifically to the one program and do not provide a foundation on which to build a variety of distributed programs (e.g., distributed operating systems, electronic mail systems, computer conferencing systems, etc.).

While conventional transport services can be used as the basis for building distributed programs, these services exhibit many organizational problems, such as the use of different data types in different machines, lack of facilities for synchronization, and no provision for a simple programming paradigm.

Distributed systems usually contain a number of different types of machines interconnected by communications networks. Each machine has its own internal data types, its own address alignment rules, and its own operating system. This heterogeneity causes problems when building distributed systems. As a result, program developers must include in

programs developed for such heterogeneous distributed systems the capability of dealing with ensuring that information is handled and interpreted consistently in different machines.

However, one simplification is afforded by noting that a large proportion of programs use a request and response interaction between processes where the initiator (i.e., program initiating a communication) is blocked out until the response is returned and is thus idle during this time. This can be modeled by a procedure call mechanism between processes. One such mechanism is referred to as the remote procedure call (RPC).

RPC is a mechanism for providing synchronized communication between two processes (e.g., program, applet, etc.) running on the same machine or different machines. In a simple case, one process, e.g., a client program, sends a message to another process, e.g., a server program. In this case, it is not necessary for the processes to be synchronized either when the message is sent or received. It is possible for the client program to transmit the message and then begin a new activity, or for the server program's environment to buffer the incoming message until the server program is ready to process a new message.

RPC, however, imposes constraints on synchronism because it closely models the local procedure call, which requires passing parameters in one direction, blocking the calling process (i.e., the client program) until the called procedure of the server program is complete, and then returning a response. RPC thus involves two message transfers, and the synchronization of the two processes for the duration of the call.

The RPC mechanism is usually implemented in two processing parts using the local procedure call paradigm, one part being on the client side and the other part being on the server side. Both of these parts will be described below with reference to FIG. 1.

FIG. 1 is a diagram illustrating the flow of call information using an RPC mechanism. As shown in FIG. 1, a client program 100 issues a call (step 102). The RPC mechanism 101 then packs the call as arguments of a call packet (step 103), which the RPC mechanism 101 then transmits to a server program 109 (step 104). The call packet also contains information to identify the client program 100 that first sent the call. After the call packet is transmitted (step 104), the RPC mechanism 101 enters a wait state during which it waits for a response from the server program 109.

The RPC mechanism 108 for the server program 109 (which may be the same RPC mechanism as the RPC mechanism 101 when the server program 109 is on the same platform

as the client program 100) receives the call packet (step 110), unpacks the arguments of the call from the call packet (step 111), identifies, using the call information, the server program 109 to which the call was addressed, and provides the call arguments to the server program 109.

The server program receives the call (step 112), processes the call by invoking the appropriate procedure (step 115), and returns a response to the RPC mechanism 108 (step 116). The RPC mechanism 108 then packs the response in a response packet (step 114) and transmits it to the client program 100 (step 113).

Receiving the response packet (step 107) triggers the RPC mechanism 101 to exit the wait state and unpack the response from the response packet (step 106). RPC 101 then provides the response to the client program 100 in response to the call (step 105). This is the process flow of the typical RPC mechanism modeled after the local procedure call paradigm. Since the RPC mechanism uses the local procedure call paradigm, the client program 100 is blocked at the call until a response is received. Thus, the client program 100 does not continue with its own processing after sending the call; rather, it waits for a response from the server program 109.

The Java™ programming language is an object-oriented programming language that is typically compiled into a platform-independent format, using a bytecode instruction set, which can be executed on any platform supporting the Java virtual machine (JVM). This language is described, for example, in a text entitled "The Java Language Specification" by James Gosling, Bill Joy, and Guy Steele, Addison-Wesley, 1996, which is incorporated herein by reference. The JVM is described, for example, in a text entitled "The Java Virtual Machine Specification," by Tim Lindholm and Frank Yellin, Addison Wesley, 1996, which is incorporated herein by reference. Java and Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Because the JVM may be implemented on any type of platform, implementing distributed programs using the JVM significantly reduces the difficulties associated with developing programs for heterogenous distributed systems. Moreover, the JVM uses a Java remote method invocation system (RMI) that enables communication among programs of the system. RMI is explained in, for example, the following document, which is incorporated herein by reference: Remote Method Invocation Specification, Sun Microsystems, Inc.

(1997), which is available via universal resource locator (URL)

<http://www.javasoft.com/products/jdk/1.1/docs/guide/rmi/spec/rmiTOC.doc.html>.

FIG. 2 is a diagram illustrating the flow of objects in an object-oriented distributed system 200 including machines 201 and 202 for transmitting and receiving method invocations using the JVM. In system 200, machine 201 uses RMI 205 for responding to a call for object 203 by converting the object into a byte stream 207 including an identification of the type of object transmitted and data constituting the object. While machine 201 is responding to the call for object 203, a process running on the same or another machine in system 200 may continue operation without waiting for a response to its request.

Machine 202 receives the byte stream 207. Using RMI 206, machine 202 automatically converts it into the corresponding object 204, which is a copy of object 203 and which makes the object available for use by a program executing on machine 202. Machine 202 may also transmit the object to another machine by first converting the object into a byte stream and then sending it to the third machine, which also automatically converts the byte stream into the corresponding object.

The communication between these machines sometimes involves, for example, repeated calls for the same information. These calls are made to a local proxy, which acts as a surrogate for the remote object in the address space of the client. Such a proxy will service the call by making a network request to the server object. Repeated calls to the same server object through a proxy can generate considerable network traffic, increasing the time and expense of obtaining the information. Accordingly, a need exists for a technique that reduces the amount of network communication in, for example, such a case.

SUMMARY OF THE INVENTION

A method consistent with the present invention transmits a request for a particular object. A response to the request is received, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the object, local to the requesting object, using the representation.

Another method consistent with the present invention receives at a machine a request for a particular object. A response to the request is transmitted, the response including first

code for constructing a representation of the object and including an indication of second code for processing, such that the construction creates an object for processing calls to the object, local to the requesting object, using the representation.

An apparatus consistent with the present invention transmits a request for a particular object. The apparatus receives a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the object, local to the requesting object, using the representation.

Another apparatus consistent with the present invention receives at a machine a request for a particular object. The apparatus transmits a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing, such that the construction creates an object for processing calls to the object, local to the requesting object, using the representation.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings are incorporated in and constitute a part of this specification and, together with the description, explain the advantages and principles of the invention. In the drawings,

FIG. 1 is a diagram illustrating the flow of call information using an RPC mechanism;

FIG. 2 is a diagram illustrating the transmission of objects in an object-oriented distributed system;

FIG. 3 is a diagram of an exemplary distributed processing system that can be used in an implementation consistent with the present invention;

FIG. 4 is a diagram of an exemplary distributed system infrastructure;

FIG. 5 is a diagram of a computer in a distributed system infrastructure shown in FIG. 4;

FIG. 6 is a block diagram of a distributed network for use in downloading smart proxies;

FIG. 7 is a flow chart of a process for downloading smart proxies within, for example, the distributed network shown in FIG. 6; and

FIG. 8 is a flow chart of a process for changing the processing performed by a smart proxy.

DETAILED DESCRIPTION

Overview

Instead of receiving a proxy that only makes network requests to the object for which it is a surrogate, a machine in a distributed system receives a smart proxy. Such a proxy can respond to calls on the object for which it is a surrogate without making any network calls to increase program efficiency, or perform processing before making a network call or after the completion of the network call to increase program functionality. The term proxy generally refers to code or other mechanism used to act as a surrogate for a remote object in the address space of a machine.

Systems transferring stubs and associated smart proxies may use a variant of an RPC or RMI, passing arguments and return values from one process to another process each of which may be on different machines. The term "machine" is used in this context to refer to a physical machine or a virtual machine. Multiple virtual machines may exist on the same physical machine. Examples of RPC systems include distributed computed environment (DCE) RPC and Microsoft distributed common object model (DCOM) RPC. A memory stores the stub and associated smart proxy, and this memory may include secondary sources such as a disk or receiving objects from the Internet.

Distributed Processing System

FIG. 3 illustrates an exemplary distributed processing system 300 which can be used in an implementation consistent with the present invention. In FIG. 3, distributed processing system 300 contains three independent and heterogeneous platforms 301, 302, and 303 connected in a network configuration represented by network cloud 319. The composition and protocol of the network configuration represented by cloud 319 is not important as long as it allows for communication of the information between platforms 301, 302 and 303. In addition, the use of just three platforms is merely for illustration and does not limit an implementation consistent with the present invention to the use of a particular number of platforms. Further, the specific network architecture is not crucial to embodiments consistent with this invention. For example, another network architecture that could be used in an implementation consistent with this invention would employ one platform as a network controller to which all the other platforms would be connected.

In the implementation of distributed processing system 300, platforms 301, 302 and 303 each include a processor 316, 317, and 318 respectively, and a memory, 304, 305, and 306, respectively. Included within each memory 304, 305, and 306, are applications 307, 308, and 309, respectively, operating systems 310, 311, and 312, respectively, and RMI components 313, 314, and 315, respectively.

Applications 307, 308, and 309 can be applications or programs that are either previously written and modified to work with, or that are specially written to take advantage of, the services offered by an implementation consistent with the present invention. Applications 307, 308, and 309 invoke operations to be performed in accordance with an implementation consistent with this invention.

Operating systems 310, 311, and 312 are typically standard operating systems tied to the corresponding processors 316, 317, and 318, respectively. The platforms 301, 302, and 303 can be heterogenous. For example, platform 301 has an UltraSparc® microprocessor manufactured by Sun Microsystems, Inc. as processor 316 and uses a Solaris® operating system 310. Platform 302 has a MIPS microprocessor manufactured by Silicon Graphics Corp. as processor 317 and uses a Unix operating system 311. Finally, platform 303 has a Pentium microprocessor manufactured by Intel Corp. as processor 318 and uses a Microsoft Windows 95 operating system 312. An implementation consistent with the present invention is not so limited and could accommodate homogenous platforms as well.

Sun, Sun Microsystems, Solaris, Java, and the Sun Logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UltraSparc and all other SPARC trademarks are used under license and are trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Memories 304, 305, and 306 serve several functions, such as general storage for the associated platform. Another function is to store applications 307, 308, and 309, RMI components 313, 314, and 315, and operating systems 310, 311, and 312 during execution by the respective processor 316, 317, and 318. In addition, portions of memories 304, 305, and 306 may constitute shared memory available to all of the platforms 301, 302, and 303 in network 319. Note that RMI components 313, 314, and 315 operate in conjunction with a JVM, which is not shown for the purpose of simplifying the figure.

Distributed System Infrastructure

Systems and methods consistent with the present invention may also operate within a particular distributed system 400, which will be described with reference to FIGS. 4 and 5. This distributed system 400 is comprised of various components, including hardware and software, to (1) allow users of the system to share services and resources over a network of many devices; (2) provide programmers with tools and programming patterns that allow development of robust, secured distributed systems; and (3) simplify the task of administering the distributed system. To accomplish these goals, distributed system 400 utilizes the Java programming environment to allow both code and data to be moved from device to device in a seamless manner. Accordingly, distributed system 400 is layered on top of the Java programming environment and exploits the characteristics of this environment, including the security offered by it and the strong typing provided by it.

In distributed system 400 of FIGS. 4 and 5, different computers and devices are federated into what appears to the user to be a single system. By appearing as a single system, distributed system 400 provides the simplicity of access and the power of sharing that can be provided by a single system without giving up the flexibility and personalized response of a personal computer or workstation. Distributed system 400 may contain thousands of devices operated by users who are geographically disperse, but who agree on basic notions of trust, administration, and policy.

Within an exemplary distributed system are various logical groupings of services provided by one or more devices, and each such logical grouping is known as a Djinn. A "service" refers to a resource, data, or functionality that can be accessed by a user, program, device, or another service and that can be computational, storage related, communication related, or related to providing access to another user. Examples of services provided as part of a Djinn include devices, such as printers, displays, and disks; software, such as programs or utilities; information, such as databases and files; and users of the system.

Both users and devices may join a Djinn. When joining a Djinn, the user or device adds zero or more services to the Djinn and may access, subject to security constraints, any one of the services it contains. Thus, devices and users federate into a Djinn to share access to its services. The services of the Djinn appear programmatically as objects of the Java programming environment, which may include other objects, software components written in

different programming languages, or hardware devices. A service has an interface defining the operations that can be requested of that service, and the type of the service determines the interfaces that make up that service.

Distributed system 400 is comprised of computer 402, a computer 404, and a device 406 interconnected by a network 408. Device 406 may be any of a number of devices, such as a printer, fax machine, storage device, computer, or other devices. Network 408 may be a local area network, wide area network, or the Internet. Although only two computers and one device are depicted as comprising distributed system 400, one skilled in the art will appreciate that distributed system 400 may include additional computers or devices.

FIG. 5 depicts computer 402 in greater detail to show a number of the software components of distributed system 400. One skilled in the art will appreciate that computer 404 or device 406 may be similarly configured. Computer 402 includes a memory 502, a secondary storage device 504, a central processing unit (CPU) 506, an input device 508, and a video display 510. Memory 502 includes a lookup service 512, a discovery server 514, and a Java runtime system 516. The Java runtime system 516 includes the Java RMI system 518 and a JVM 520. Secondary storage device 504 includes a Java space 522.

As mentioned above, distributed system 400 is based on the Java programming environment and thus makes use of the Java runtime system 516. The Java runtime system 516 includes the Java API libraries, allowing programs running on top of the Java runtime system to access, in a platform-independent manner, various system functions, including windowing capabilities and networking capabilities of the host operating system. Since the Java API libraries provides a single common API across all operating systems to which the Java runtime system is ported, the programs running on top of a Java runtime system run in a platform-independent manner, regardless of the operating system or hardware configuration of the host platform. The Java runtime system 516 is provided as part of the Java software development kit available from Sun Microsystems, Inc. of Mountain View, CA.

JVM 520 also facilitates platform independence. JVM 520 acts like an abstract computing machine, receiving instructions from programs in the form of bytecodes and interpreting these bytecodes by dynamically converting them into a form for execution, such as object code, and executing them. RMI 518 facilitates remote method invocation by allowing objects executing on one computer or device to invoke methods of an object on

another computer or device. Both RMI and the JVM are also provided as part of the Java software development kit.

Lookup service 512 defines the services that are available for a particular Djinn. That is, there may be more than one Djinn and, consequently, more than one lookup service within distributed system 400. Lookup service 512 contains one object for each service within the Djinn, and each object contains various methods that facilitate access to the corresponding service. Lookup service 512 is described in U.S. patent application entitled "Method and System for Facilitating Access to a Lookup Service," which was previously incorporated herein by reference.

Discovery server 514 detects when a new device is added to distributed system 400, during a process known as boot and join (or discovery), and when such a new device is detected, the discovery server passes a reference to lookup service 512 to the new device so that the new device may register its services with the lookup service and become a member of the Djinn. After registration, the new device becomes a member of the Djinn, and as a result, it may access all the services contained in lookup service 512. The process of boot and join is described in U.S. patent application entitled "Apparatus and Method for providing Downloadable Code for Use in Communicating with a Device in a Distributed System," which was previously incorporated herein by reference.

A Java space 522 is an object repository used by programs within distributed system 400 to store objects. Programs use a Java space 522 to store objects persistently as well as to make them accessible to other devices within distributed system 400. Java spaces are described in U.S. patent application serial no. 08/971,529, entitled "Database System Employing Polymorphic Entry and Entry Matching," assigned to a common assignee, and filed on November 17, 1997, which is incorporated herein by reference. One skilled in the art will appreciate that an exemplary distributed system 400 may contain many lookup services, discovery servers, and Java spaces.

Data Flow in a Distributed Processing System

FIG. 6 is a block diagram of an object-oriented distributed network 600 connecting machines 601 and 606, such as computers or virtual machines executing on one or more computers, or the machines described with reference to FIGS. 3, 4, and 5. Network 600 transmits proxies, some of which may be smart proxies. A smart proxy includes code for

performing processing associated with a call. For example, a smart proxy may perform a caching operation for read-only data for later reference. When a call is made for that data, the smart proxy may obtain it locally and provide it to a user without making another call for the data, which may occur transparent to the user. An example of such read-only data is a particular installation time. The first time a call is made for the installation time, for example, a smart proxy locally caches that value, and when a subsequent call is made for the installation time, the smart proxy locally retrieves the value.

Another example of smart proxy processing involves use of a serialized object for transmitting data to a data bank storing information. In this example, a call is made to a smart proxy, which receives an object, serializes the object on the client machine into an array of bytes, and transmits the array of bytes to a server. The server only stores the serialized object, avoiding the requirement to download code, and it provides a key for the object to the client machine. When the client machine wants to retrieve the data, the smart proxy transmits the key to the server, receives in response the serialized object, reconstructs the object, and provides it the user.

Other examples of uses of smart proxies include processing for debugging, call logging, and monitoring system performance. Another example involves the use of a smart proxy for local data verification, as explained in U.S. patent application filed on even data herewith, assigned to a common assignee, and entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," which is incorporated herein by reference. Many other uses for smart proxies are possible for performing processing associated with a call.

Network 600 includes a client machine 601 containing RMI 602 and associated code 603. A server machine 606 includes RMI 607 and remote object 608. In operation, RMI 602 transmits a call or request 609 to RMI 607, requesting a particular stub object. RMI 607 returns a response 610 including requested stub 605 embedded within a smart proxy 604. The response may be transmitted as a stream. Streams used in the Java programming language, including input and output streams, are known in the art and an explanation, which is incorporated herein by reference, appears in, for example, a text entitled "The Java Tutorial: Object-Oriented Programming for the Internet," pp. 325-53, by Mary Campione and Kathy Walrath, Addison-Wesley, 1996.

The response may include information so that client machine 601 can reconstruct the stub object in smart proxy 604. When a set of object types is limited and is the same on machines 601 and 606, a receiving machine typically requires the object's state and a description of its type because the object's code is already present on all network machines. Alternatively, machine 606 uses RMI 607 to provide more flexibility, allowing code to be moved when necessary along with information or the object's state and type. Additionally, a transmitting machine may include in the object an identification of the type of object transmitted, the data constituting the state of the object, and a network-accessible location in the form of a URL for code that is associated with the object. URLs are known in the art and an explanation, which is incorporated herein by reference, appears in, for example, a text entitled "The Java Tutorial: Object-Oriented Programming for the Internet," pp. 494-507, by Mary Campione and Kathy Walrath, Addison-Wesley, 1996.

When client machine 601 receives response 610, it identifies the type of transmitted object. Machine 601 contains its own RMI 602 and code 603 for processing of objects, and it may create stub object 605 using the object type, the state information, and code for the object. If code for the object is not resident or available on machine 601 and the stub object does not contain the code, RMI 602 may use a URL from the object to locate the code and transfer a copy of the code to client machine 601. Because the code is bytecodes and is therefore portable, client machine 601 can load the code into RMI 602 to reconstruct the object. Thus, client machine 601 can reconstruct an object of the appropriate type even if that kind of object has not been present on the machine before.

When creating stub object 605, RMI 602 does not necessarily know that the stub is itself a smart proxy 604. Smart proxy 604 may perform processing at client machine 601 before or after response 610 and may supply all processing without resorting to call 609 to the object for which the proxy acts. Therefore, smart proxy 604 may perform all processing locally when client machine 601 makes a call or request 611 to invoke a method on smart proxy 604. These proxies are downloadable by the same methods as disclosed in U.S. patent application serial no. 08/950,756, filed on October 15, 1997, and entitled "Deferred Reconstruction of Objects and Remote Loading in a Distributed System," which is incorporated herein by reference.

Transmission of Smart Proxies

FIG. 7 is a flow chart of a process 700 for downloading and using smart proxies within, for example, the distributed network shown in FIG. 6. A client machine transmits a call or request for a particular object (step 701), and a server machine receives the call (step 702). In response, the server machine returns a smart proxy with an embedded stub (step 703), and the proxy acts as a representation of the requested object. After receiving the smart proxy, the client machine invokes a method on it (step 704). According to the code within the smart proxy, the client machine containing the smart proxy determines if preprocessing is required (step 705). If so, the processing is performed locally by the client machine using the smart proxy (step 706).

The client machine then determines if the method called on the smart proxy may be serviced locally (step 707). If so, the client machine performs the local processing for the call (step 711). If not, the client machine calls the remote object (step 708). The remote processing is performed (step 709), and the result of the remote processing is returned to the client machine (step 710).

The client machine determines, according to code in the smart proxy, if post-processing as a result of the call is required (step 712). If so, it locally performs the post-processing using code in the smart proxy (step 713). The smart proxy then returns the method call result (step 714) in response to the call on the smart proxy in step 704.

FIG. 8 is a flow chart of a process 800 for changing the processing performed by a smart proxy. When processing is invoked (step 801), a client machine determines if updated processing is required (step 802). Such information may be contained within the smart proxy itself in that it may determine when or under what particular circumstances it requires updated processing code. If updated processing is required, the code for that processing is downloaded and the smart proxy is updated at the client machine to perform that processing (step 803). The smart proxy then performs at the client machine the processing according to the updated code (step 804).

Machines implementing the steps shown in FIGS. 7 and 8 may include computer processors for performing the functions, as shown in FIGS. 3, 4, 5, and 6. They may include modules or programs configured to cause the processors to perform the above functions. They may also include computer program products stored in a memory. The computer program products may include a computer-readable medium or media having computer-

readable code embodied therein for causing the machines to perform functions described above. The computer-readable media may include computer data signals embodied in a carrier wave and representing sequences of instructions which, when executed by a processor, cause the processor to securely address a peripheral device at an absolute address by performing the method described in this specification. The media may also include a data structure for use in performing the method described in this specification.

Although the illustrative embodiments of the systems consistent with the present invention are described with reference to a computer system implementing the Java programming language on the JVM specification, the invention is equally applicable to other computer systems processing code from different programming languages. Specifically, the invention may be implemented with both object-oriented and nonobject-oriented programming systems. In addition, although an embodiment consistent with the present invention has been described as operating in the Java programming environment, one skilled in the art will appreciate that the present invention can be used in other programming environments as well.

While the present invention has been described in connection with an exemplary embodiment, it will be understood that many modifications will be readily apparent to those skilled in the art, and this application is intended to cover any adaptations or variations thereof. For example, different labels or definitions for the smart proxies may be used without departing from the scope of the invention. This invention should be limited only by the claims and equivalents thereof.

WHAT IS CLAIMED IS:

1. A method for receiving objects in a distributed system comprised of multiple machines, comprising:
 - transmitting a request for a particular object; and
 - receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.
2. The method of claim 1, further including using the code to construct the representation.
3. The method of claim 2 wherein the using includes downloading code to update the processing.
4. A method for transmitting objects in a distributed system comprised of multiple machines, comprising:
 - receiving at a machine a request for a particular object; and
 - transmitting a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.
5. An apparatus for receiving objects in a distributed system comprised of multiple machines, comprising:
 - a module configured to transmit a request for a particular object; and
 - a module configured to receive a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.
6. The apparatus of claim 5, further including

a module configured to use the code to construct the representation.

7. The apparatus of claim 6 wherein the use module includes
a module configured to download code to update the processing.
8. An apparatus for transmitting objects in a distributed system comprised of multiple machines, comprising:
a module configured to receive at a machine a request for a particular object; and
a module configured to transmit a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.
9. A system for transmitting objects in a distributed system comprised of multiple machines, comprising:
a first machine;
a second machine;
a network connecting the first machine with the second machine; and
an apparatus for receiving objects, the apparatus including:
a module configured to transmit a request for a particular object; and
a module configured to receive a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.
10. The system of claim 9, further including
a module configured to use the code to construct the representation.
11. The system of claim 10 wherein the use module includes
a module configured to download code to update the processing.

12. A system for transmitting objects in a distributed system comprised of multiple machines, comprising:
- a first machine;
 - a second machine;
 - a network connecting the first machine with the second machine; and
 - an apparatus for transmitting objects, the apparatus including:
 - a module configured to receive at a machine a request for a particular object;
 - and
 - a module configured to transmit a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.
13. A computer program product, comprising:
- a computer-readable medium containing instructions for controlling a computer system to perform a method, the method including:
 - transmitting a request for a particular object; and
 - receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.
14. The computer program product of claim 13, further including using the code to construct the representation.
15. The computer program product of claim 14 wherein the using includes downloading code to update the processing.
16. A computer program product, comprising:
- a computer-readable medium containing instructions for controlling a computer system to perform a method, the method including:

receiving at a machine a request for a particular object; and
transmitting a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.

17. An article of manufacture specifying a representation of an object stored in a computer-readable storage medium and capable of electronic transmission between machines in a distributed system, the article of manufacture comprising:

an object to be transmitted from a first machine to a second machine in a response to a request, the object including an indication of code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.

18. An apparatus for receiving objects in a distributed system comprised of multiple machines, comprising:

means for transmitting a request for a particular object; and

means for receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to a requesting object, using the representation.




FIG. 1




FIG. 2

3 / 8




FIG. 3

4 / 8




FIG. 4

5 / 8




FIG. 5


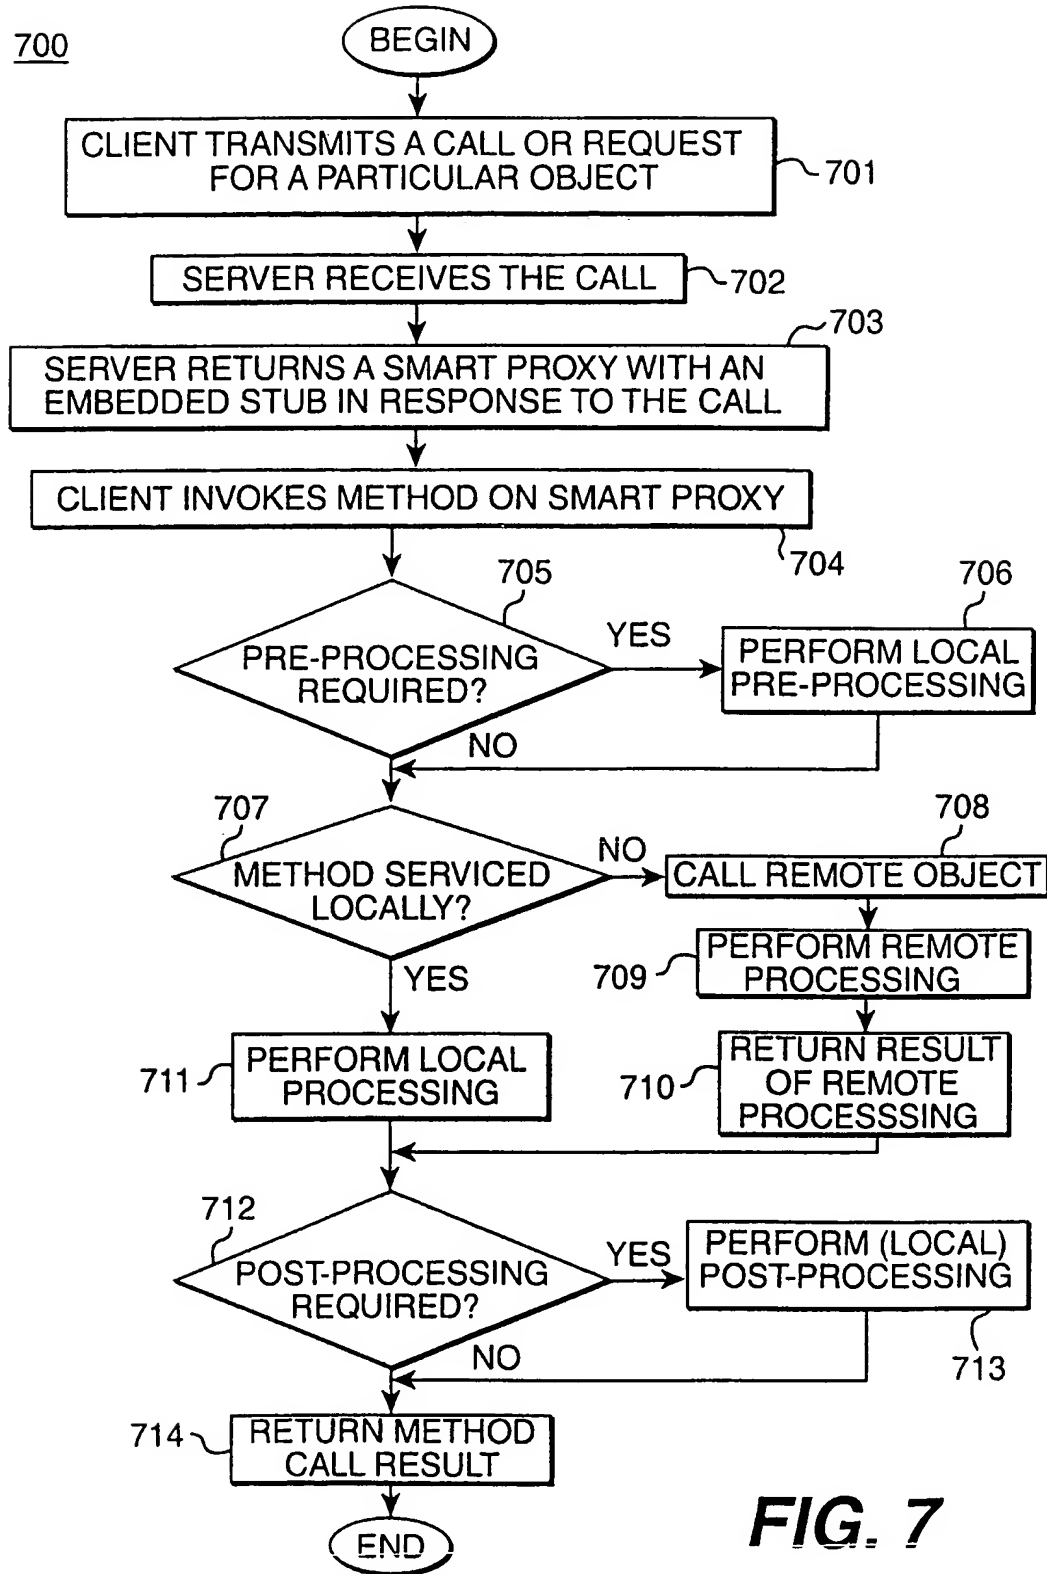


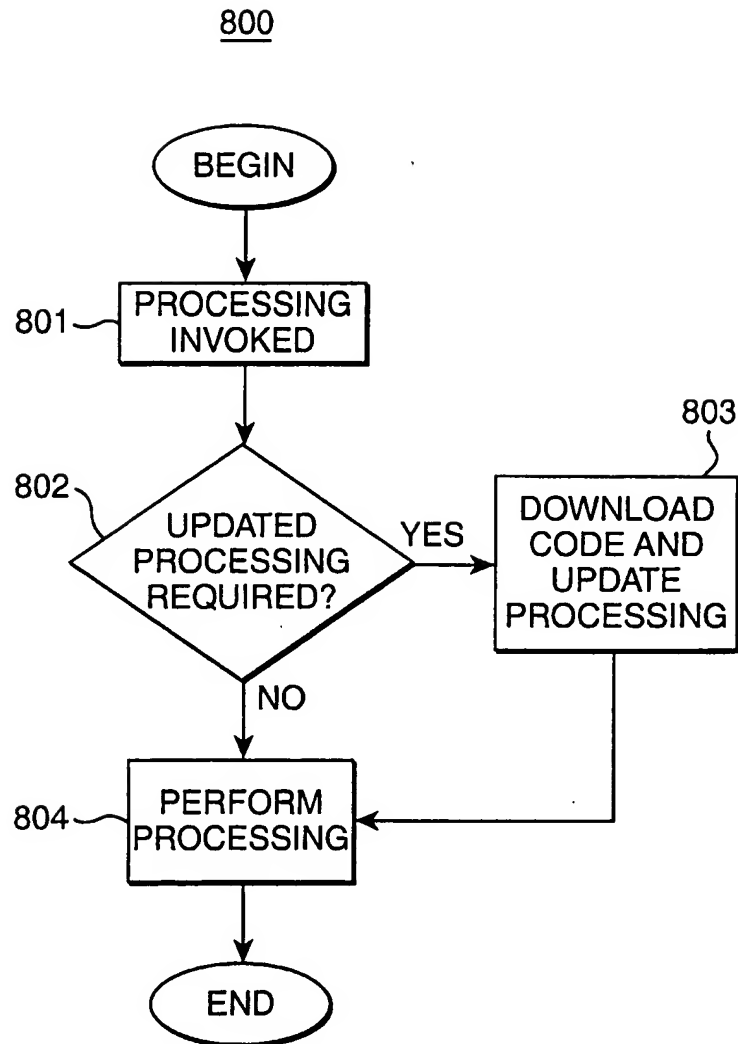
FIG. 6

7 / 8

700

**FIG. 7**

8 / 8

**FIG. 8**

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 99/03221

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	IBM: "Chapter 6 - Distributed SOM (DSOM)" SOMOBJECTS DEVELOPER TOOLKIT USERS GUIDE, VERSION 2.1, October 1994 (1994-10), pages 6-1-6-90, XP002047926 US ISSN: 0003-6935	1,2,4-6, 8-10, 12-14, 16-18
A	page 6-66, line 46 - page 6-67, line 5 page 6-81, line 25 - page 6-84, line 31 ---	3,7,11, 15
X	ROBERT ORFALI, DAN HARKEY & JERI EDWARDS: "The Essential Distributed Objects Survival Guide" 1996, JOHN WILEY & SONS, INC., US XP002108734	1,2,4-6, 8-10, 12-14, 16-18
A	page 206, line 26 - line 35 --- -/--	3,7,11, 15

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

9 July 1999

Date of mailing of the international search report

22/07/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Fonderson, A